UCRL-ID-132280

Using Artificial Neural Networks and the Genetic Algorithm to Optimize Well-Field Design: Phase I Final Report

V.M. Johnson
L.L. Rogers

March 1998

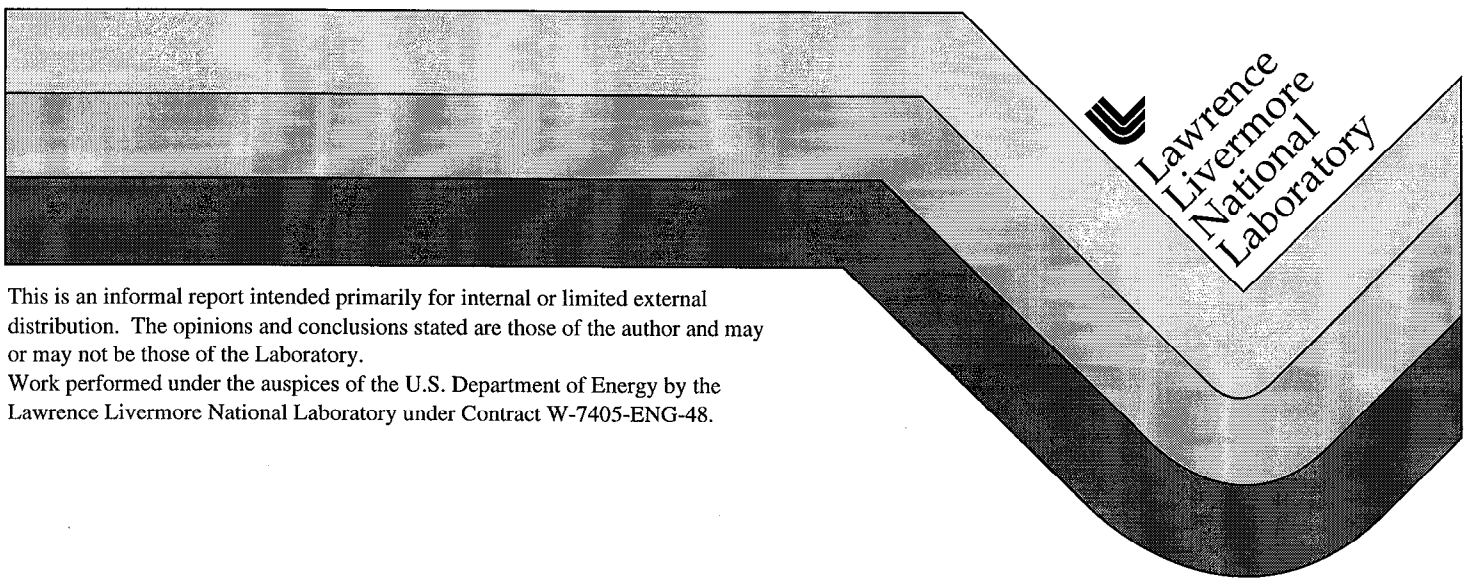

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

Using Artificial Neural Networks and the Genetic Algorithm to Optimize Well-Field Design:

Phase I Final Report

October 15, 1998

Virginia M. Johnson and Leah L. Rogers
Lawrence Livermore National Laboratory
Livermore, California

Table of Contents

Introduction

Statement of the Problem. Reservoir simulation is a well-established component of reservoir management throughout much of the petroleum industry. Black oil simulators and more complex compositional, thermal, and chemical models are used as forecasting tools in both the day-to-day operational management of production facilities and longer-term field development planning. As yet, however, little use has been made of reservoir simulation coupled with systematic optimization techniques. The main advantage of applying these mathematical tools to decision-making problems is that they are less restricted by human imagination than conventional case-by-case comparisons. As the number of competing engineering, economic, and environmental planning objectives and constraints increases, it becomes difficult for human planners to track complex interactions and select a manageable set of promising scenarios for examination. Using optimization techniques, the search can range over all possible combinations of variables, locating strategies whose effectiveness is not always obvious to planners. Optimization also generates large sets of promising scenarios from which planners can choose.

The single biggest obstacle to the application of optimization techniques using a reservoir simulator as the forecasting tool is the computational time required to complete a single simulation. Even the examination of 10 variations on a well-field design becomes cumbersome when a single run requires hours or days to complete. Extending the use of these simulators into optimization schemes involving hundreds or thousands of runs poses a computational problem bigger than most organizations are willing or able to tackle. The ANN-GA/SA solution to this problem is to train artificial neural networks (ANNs) to predict selected information that the simulator would normally predict. A heuristic search engine, either the genetic algorithm (GA) or simulated annealing (SA), searches for increasingly better strategies (such as the most productive in-fill drilling pattern or the best distribution of steam injection wells), using the trained networks to evaluate the effectiveness of each strategy in place of the original simulator. This substitution has been shown to reduce the time needed to evaluate pump-and-treat groundwater remediation strategies by a factor of nearly a million, enabling the evaluation of millions of strategies in a matter of days on conventional workstations. After analysis of the results of the search, the best-performing strategies are submitted to the original simulator to confirm their performance.

This report describes a one-year pilot project to assess the applicability of the ANN-GA/SA approach to the proposed water flood of a deep water reservoir that has been in production for 2.5 years. The management goal was to identify the best set of injection well locations to maximize some economic measure of performance over a seven-year planning horizon. The ANN-GA/SA methodology was originally developed for use on 2-D groundwater remediation problems. The thrust, then, of the pilot project was to determine how well these methods would translate to 3-D simulation of multiphase flow, with minimal adaptation of procedures.

ANN-GA/SA Approach to Optimization. Simulation-optimization, a term which refers to the coupling of models to optimization drivers, has received extensive attention in the groundwater remediation literature. The goal of optimization for this type of problem is usually to find one or more combinations of extraction and injection well locations that will at least contain and preferrably clean up the contamination at minimum cost or time. Although the number of well combinations is potentially infinite, it has been customary in groundwater optimization work to prespecify a grid of potentially good well locations and then formulate the search to locate the most time- or cost-effective subset of those locations which meets remediation goals.

Early optimization work at the western U.S. Superfund site at which the ANN-GA/SA methodology was developed used 20 preselected extraction locations with fixed pumping rates and searched for the subset producing the smallest volume of treated water (a convenient surrogate for cost) which contained the contamination over a 40-year planning period [1]. Later work focused on 28 fixed-rate extraction and injection locations in a multiple-objective search which balanced cost-efficiency with mass-extraction performance, while meeting a containment constraint over a 50-year planning period [2].

Regardless of the problem formulation or the type of search technique employed, key components of the cost function for a particular well pattern are evaluated by a contaminant transport model which assesses the impact of the well pattern on the distribution of the contamination over some period of time. Even in 2-D, numerical models of this sort can take hours to evaluate a single pattern on a conventional workstation. As the resolution, dimensionality, and heterogeneity of the models increase, the time required for this evaluation can extend to days. Since even the most efficient, deterministic search techniques usually need to evaluate hundreds of patterns, the modeling step becomes a major computational bottleneck in the optimization of realistic environmental engineering problems. Much of the work in this area has accepted the modeling bottleneck as a given, sometimes simplifying the situation by analyzing smaller-scale problems or using simpler models (see [3] for several examples) or seeking to reduce the number of times the model must be called by increasing the efficiency of the search itself [4]. Other work has highlighted the acute need to break the bottleneck by suggesting that multiple realizations of geologic parameters are necessary to address the uncertainty in the simulation models [5, 6].

Work intended to confront the modeling bottleneck directly falls into one of two camps. The first approach involves reducing the execution time required by the model through parallel algorithms and computer architectures [7, 8]. This represents a "rich man's" approach because of the costs normally associated with gaining access to computer resources of this kind. The ANN-GA/SA approach, in contrast, confronts the problem by training neural networks to predict selected model results. The trained networks, rather than the original model, are then used by a some seasrch technique to obtain performance predictions in fractions of a second.

The network architecture used for this prediction task is a multilayer perceptron, trained by the standard backpropagation learning algorithm [9]. Training and testing examples are obtained by associating well pattern variations with selected outcomes such as the amount of contamination that has been removed, the highest remaining concentrations after treatment is complete, and whether or not contamination has spread beyond certain boundaries. The examples are drawn from a knowledge base initially created by running the contaminant transport model on a *representative sample* of well patterns. Since there are no dependencies among the model runs, they can be distributed over a network of processors using only the basic remote file system and execution facilities that are now a standard part of most network environments.

Although the trained nets can be coupled with a variety of search techniques, heuristic search techniques (namely, the genetic algorithm [10] and simulated annealing [11]) have been the methods of choice for three reasons. First, they are probabilistic, rather than deterministic, search techniques that have been developed by analogy to natural processes. Second, since they both employ direct function evaluation rather than derivatives of functions, they allow more complex integration of different components. In other words, both apples and oranges can be optimized simultaneously. Finally, the heuristic methods represent a philosophy of search that is especially well suited to design optimization problems [12]. The contaminant transport models used to evaluate the effectiveness of each well pattern are crude approximations of reality. Their utility lies more in outlining broad hydrological design principles applicable to a given site than in predicting precise outcomes. Furthermore, there are many practical engineering, managerial, and political constraints that cannot easily be quantified in a cost function. Consequently, employing a search strategy oriented toward producing one or a handful of best solutions is not likely to be well-received by engineers and planners. Instead, a search technique generating a wide range of potentially useful solutions, which are subsequently analyzed for their common properties, is generally more useful. Designers can then select especially interesting solutions to incorporate into their detailed designs or simply follow the general principles suggested by the analyses.

The main components of the ANN-GA/SA methodology are shown in Fig. 1. It is important to note that, since the set of optimal solutions generated by the search engines is obtained by an ANN estimation process that introduces a certain degree of error, the final step in the methodology is to submit that optimal set to the original simulator for verification. The updated performance measures supplied by the simulator on this manageable set of scenarios are the ones which are used in subsequent decision-making.

Design Optimization in Petroleum Engineering. Reservoir simulation is now a well-established component of reservoir management, as indicated by the role it is given in both general discussions [13, 14] and case studies [15, 16] of reservoir management practices. But the

use of these simulators in a structured, formal search for more effective recovery strategies is still in its infancy. Typically, the simulator is used to evaluate a small set of development scenarios that have been selected to test specific hypotheses (see, for example, [17-19]). A few researchers, however, have examined more structured approaches.

Aanonsen et al. [20] apply concepts from experimental design and response surfaces to optimize a reservoir response variable (such as oil production rate) according to reservoir management parameters (such as well location and flow rates). Their largest example involved 240 one-hour runs of a 5500 grid block 3-D model of a fluvial reservoir. The goal was to build a response surface of discounted oil production from sample inputs consisting of the x and y coordinates of a single producer and the x coordinate of a single injector. To account for uncertainties in the flow field, these three inputs were crossed, as in an experimental design, with eight different realizations of the deposition of channel sands. The response surface was examined for distinct maxima, which became the optimal solutions to the problem. This work is similar to the ANN-GA/SA methodology in that the results of a sample of simulations are used to build surfaces which are then searched for solutions. In the ANN-GA/SA approach, however, the sampling is performed to create a re-usable archive of data. The archive provides the examples from which many different networks figuring in many different searches are drawn.

Wackowski et al. [21] employ decision analysis techniques to examine over 2500 expansion, investment, operational, and CO_2 purchase/recompression scenarios to maximize net present value of a project at the Rangely Weber Sand Unit. This ambitious, long-range project pulled together information from many sources (including expert opinion, economic spreadsheet models and reservoir models) into decision trees, from which the highest probability paths were selected. The reservoir model combined the vertical response of a single detailed cross-section with the areal response of a full-field streamtube model to obtain full-field forecasts of injected and produced fluids. Since several techniques were used to reduce the number of paths in the decision tree which required full examination, it is unclear how many scenarios the simulator actually evaluated. This approach to optimization is similar to the ANN-GA/SA methodology in that they both examine very large numbers of alternatives. The techniques, however, are quite dissimilar in their identification of optimal solutions. Unless it is exhaustive of all possibilities, which is unlikely in a real-world problem, a decision tree can only select solutions from paths that have been anticipated by its designers. Optimization techniques, in contrast, can uncover combinations of inputs which produce results which were not anticipated.

The most classic application of optimization techniques to facility design is given by Fujii and Horne [22]. They compare three different search techniques (a derivative-based method, the polytope method, and the GA) as applied to the optimization of a networked production system by varying parameters such as separator pressure, diameters of tubing, pipeline vs surface choke, and so on. Calculations were restricted to relatively simple production rate equations because the

use of a reservoir simulator was judged to be too time-consuming. Later, Bittencourt and Horne [23] used a GA combined with economics and simulation to determine the optimal relocation of wells in a proposed 33-well layout and the best platform location. Their experiences reinforce the motivation behind the ANN-GA/SA approach: that the advantages of optimization techniques will not be fully exploited until some method is found to reduce the computational burden imposed by the reservoir simulator.

These recent advances suggest that the petroleum engineering field is beginning to pay attention to more structured approaches to the optimization of development strategies. Use of the ANN-GA/SA approach can promote further interest in this process by alleviating the computational bottleneck created by the reservoir simulator. This is accomplished not by eliminating the simulator from the optimization loop, as is done in the Fujii and Horne work, but by capturing simulator predictions in the weights of an artificial neural network. In this way, the results of the optimization continue to benefit from the increased accuracy of predictions that a reservoir simulator can provide without having to pay the full price in computational time. The critical role played by the simulator is reinforced when the best-performing scenarios generated by the search are submitted to it for validation.

Reservoir Description

The Pompano Field in the deep water Gulf of Mexico is the test site for this project. BP and Kerr-McGee are joint operators of this field, which has been in production since April, 1995. They have developed and callibrated a reservoir model, using Landmark's VIP® simulator, for the Miocene section.

Field. The Pompano field consists of multiple turbidite reservoirs in a variety of structural traps and settings. An intrusive salt body and a large counter-regional growth fault are important structural features in the field. The field is divided into three areas, shown in Fig. 2. To the north and northwest of the fault is the downthrown Pliocene which consists of 10 independent, stacked reservoirs. It is generally underlain by the more sheet-like part of the salt body. The upthrown Pliocene is south of the salt and growth fault. Its reservoirs are a group of related channel sand deposits. An older Miocene channel complex lies to the southeast, in the syncline between the Pompano and Mickey salt bodies. This Miocene complex consists of an interconnected group of turbidite sands. More than two thirds of the total recoverable reserves are estimated to be in the Miocene portion of the field; and our application will focus on this area as the reservoir to be managed.

Reservoir. The Miocene reservoir sands were deposited as mid-slope turbidites in a large, aggradational channel complex. There is significant connectivity between channels as younger channels frequently eroded into previously deposited ones. Pressure depletion in successively

drilled wells suggests that most of the reservoir complex is in pressure and fluid continuity. Grain size ranges from very fine to medium, with the bulk being fine grained. The average thickness of the Miocene sand is 50' net feet of oil (NFO) in a vertical target interval of 300' to 400', and the thickest sand penetrated is 110' NFO in a single sand.

Trap. A north-south trending channel system draped over an east-west trending structural nose forms the trap. The channel sands are laterally confined by the shales and silty shales of the overbank deposits. An oil-water contact at -10,200' true vertical depth sub-sea (TVDSS) has been drilled on the southern edge of the field and is implicated on the north/northwest end by seismic interpretation and water production. Maximum hydrocarbon column height is approximately 600'. The large aquifer system below, estimated to be three-fold larger than oil-in-place, is judged to be an advantage to help offset pressure losses during reservoir depletion.

Production. The Miocene oil has very favorable properties which help in achieving high production rates. API gravity is 32°, viscosity is 0.38 cp, and the gas-oil ratio (GOR) was initially 1037 scf/stbo and is climbing with increased production. The very restricted range of variability in the producing wells emphasizes the connectivity in the Miocene reservoirs. There are 12 production wells in operation, five drilled from the platform to the north during Phase I, and seven drilled from a subsea template to the south during Phase II (see Fig. 3). The average initial flow rate was 788 stb/d for the five Phase I wells and 6343 stb/d for the seven Phase II wells. The gas and oil production decline curves for 2.67 years of production (from April, 1995 through December, 1997) and seven additional years of simulated production under injection are shown in Fig. 4. The cap on gas production is a function of surface facility limitations.

Simulator. The heterogeneous anticlinal turbidite reservoir was discretized first into an approximately three million cell block model at seismic resolution. It was then scaled up to a 40,000 cell block simulation model, implemented in VIP®, with dimensions of 40 x 40 in plan view and 25 layers. At the time the project was initiated, only a two-cpu software license was available for the simulator. So, the simulation time-frame, which would normally have been 15 - 20 years, was cut to seven years to enable the knowledge base creation phase to be completed with dispatch. Seven-year simulations of the existing producers plus one to four injectors required an average of 3.5 hours to complete on a dedicated Sun UltraSparc 2 workstation.

Management Question

The planning question posed by BP is whether a water injection program will improve production from the Miocene. We created a candidate pool of 25 locations for injector wells based on high transmissivity, spatial coverage, and economics (see the methodology section below for details). The optimization problem was then formulated to search for the

combinations of one to four injector wells which maximize simple net profit, subject to facility constraints. Although only one or two injectors were being considered by the Pompano asset team at the time this project was initiated, the problem scope was expanded to include the possiblity of a more aggressive program. Management time horizons of both three and seven years were examined.

Assumptions. For this management formulation, the following assumptions were made:

1) The maximum time-period over which alternative water flood scenarios would be assessed is January 1, 1998 to January 1, 2005.

2) A candidate pool of 25 injection sites, including both existing production wells and newly-drilled injectors, would be developed. Given this candidate pool, the search would identify the particular subsets, which could vary in size from one to four wells, which maximize some measure of economic performance.

3) The Phase I producers located in the northern portion of the field would be considered for conversion to injectors; but, for engineering reasons, the Phase II producers would not.

4) Due to limitations on unused slots, only two new injectors could be drilled from the northern platform. Drilling any injectors in the southern section would require the emplacement of a new subsea template, from which up to four injectors could be drilled.

5) The 12 production wells would continue to operate as they do now, except for any that were converted to injectors.

6) The precise implementation of the water flood would be kept simple. On January 1, 1998, all sites in the well combination would commence injection and would continue to do so for the duration of simulation. Individual flow rates would be capped at 20,000 bbl/day; but actual rates would be determined by the reservoir simulator's own algorithms. No phasing of either injection or production wells would be considered.

7) Existing surface facilities constraints would be maintained. However, if a well combination's total demand for sea water to inject were to exceed the current limit of 40,000 bbl/day, upgrading would be permitted.

Cost Estimates. The estimates for costs are as follows:

1) Conversion of producers: Conversion of a Phase I producer into an injector is estimated at $7 million for the first injector and $3 million for each subsequent injector.

2) New injector wells: New locations are considered in two cost categories. A north-south dividing line is drawn to separate wells which can be drilled from the platform and those which must be drilled from the new subsea template in the south (see Fig. 3). Locations within reach of the platform can be drilled for $13 million each. Locations in the southern portion of the field require an up-front investment of $25 million to move the drilling platform into place and install the template. Each well would then cost $13 million to drill.

3) Seawater pumping facilities: Well combinations whose combined peak injection rates exceed 40,000 bbl/day will necessitate upgrading the facilities at an estimated cost of $2 million for each additional 30,000 barrels pumped.

4) Maintenance and operation (M&O): M&O costs associated with the 12 existing producers are estimated at $182k/well/year. M&O costs for injectors are estimated at $1 million/year for the first injector and $500k/year for each additional injector.

5) Value of produced oil/gas: The oil price used in the net profit calculations was $15.50/bbl. The gas price was $2.50/mcf. These values were based on New York Mercantile Exchange quotes from May, 1998.

Performance Measure (Objective Function). The measure used to evaluate the performance of individual well combinations and serve as the objective function to be optimized is simple net profit (SNP). This measure is the sum of all revenues from sale of the produced oil and gas over the time period being evaluated minus the sum of the capital and M&O costs detailed above for the same period. No discounting or inflation factors were taken into account. By using a simple measure such as this, the number of individual estimates of oil and gas production required for the calculations could be kept to a minimum: one estimate each of cumulative oil and cumulative gas production over the desired time-frame. For this particular problem, nothing was lost by optimizing on the basis of the simpler formulation. For the 550 well combinations comprising the knowledge base, the squared correlation over seven years between SNP and net present value (which was calculated using a 0.10 discount factor and a 0.03 inflation factor) was $r^2 = 0.99$.

Application of the ANN-GA/SA Methodology

Fig. 1 shows the general flow of the methodology. The application of each component to the Pompano water flood problem is discussed below.

Create a Knowledge Base of Simulations. This is the most critical component in the entire process and consists of several steps:

Define the Problem Scope - In this step, the boundaries of the problem to be optimized are determined. The decisions made in this step will guide the sampling of representative runs for the reservoir simulations and, as a result, will set the limits within which management questions can be asked. Most of the decisions that are made at this time are embodied in the assumptions detailed above. One critical issue is the maximum time-frame over which performance will be evaluated. The maximum time-frame, seven years in this case, determines the simulation period for the reservoir simulation runs. By saving intermediate yearly results, this time-frame can be shortened, if desired. But it cannot be extended without further simulation. For the Pompano

problem, the seven-year time-frame was selected for practical reasons concerning software licensing limitations.

Another set of issues involves separating factors in the problem which will be held constant from those that will be allowed to vary (i.e. the "decision variables", in decision theory terminology). For example, one of the assumptions listed earlier is that production at the existing wells will continue as before, except for any that are converted to injectors. This means that field development scenarios that involve the drilling of additional production wells cannot be considered later on because that option will not have been included in the sampling plan from which the knowledge base is created. For the Pompano problem, the only variables are 1) the size of the well combinations (from one to four) and 2) which particular wells, from the candidate pool of 25, will compose the combination.

Finally, it is necessary to identify the output variables that will go into the calculation of objective functions, such as the SNP measure defined earlier. At this stage, it is most important to define the performance measures (e.g. gas/oil production) and parameters (e.g. water injection volumes) that must be calculated by the simulator, since these decisions will determine the type and timing of output saved from each run. VIP® provides a wealth of information at each time step, ranging from production figures at the well-, region-, and field-level to updated 40,000-cell arrays of pertinent physical properties. While all that information can be archived for later exploitation, only information pertinent to the management questions likely to be posed needs to be saved. For the Pompano problem, it was anticipated that only production-related objective functions would be of interest; so, no spatial information such as the distribution of pressures or oil-in-place was archived.

Select the Candidate Pool of Well Locations - In theory, injection could occur at any of the 40,000 cell blocks comprising the reservoir model. In practice, there will be geological and engineering constraints on the siting and completion of wells. Furthermore, it is desirable to restrict consideration to some manageable number of locations, to avoid wasting simulation and search time on unprofitable scenarios. For the Pompano problem, this manageable number was set at 25, largely based on past experience with the groundwater examples cited in the literature review. The production criteria described below should be considered only suggestive of those which could be applied.

The initial candidate pool consisted of all five Phase I producers, included because conversion is less expensive than drilling a new well and because more is known about the reservoir at those points, and 50 new locations. The new locations were selected as follows: Each of the 1600 (40 x 40 in plan view) columns in the model grid was examined to locate those columns with five or more (of the 25 possible) layers having either an x- or y-transmissivity greater than 1.0. A 10 x 22 block in the southeastern corner of the grid was removed from consideration because the high transmissivities in that area were due to intersection with the aquifer. Of the 302 columns

meeting these criteria, 50 were chosen, randomly but with some manual adjustment to improve spatial dispersion, for evaluation. All 55 initial candidates were submitted to the the simulator as 1-well injection scenarios, ranked by the total hydrocarbon production (i.e. oil plus gas in oil-equivalent units) after seven years of injection, and compared to the no-injection baseline production case. The final 25 locations shown in Fig. 4 consist of the top-ranked 21 locations and four of the five Phase I wells that at least performed better than the baseline case.

Sample over the Decision Variables - This process begins by setting an overall target size for the knowledge base, 550 in this case, and sampling over the decision variables until that size is achieved. There is an approximate relationship between the number of decision variables and the number of examples required for ANN training and testing; but this relationship is also affected by the complexity of the physical relationships being modeled by the ANNs. An earlier 2-D groundwater remediation problem having 30 prospective well locations had successfully employed a total knowledge base of 400 examples (300 for training and 100 for testing). For the 3-D Pompano problem, targets of 400 training and 150 testing examples were set. The adequacy of these targets will be discussed in later sections.

The examples in the knowledge base set aside for ANN training contained the no-injection baseline case and all 25 1-well injection combinations. The remaining 374 training examples were generated in a three-step process:

1) randomly select the size, from 2-4, of the combination,
2) randomly select specific well locations, from the set of 25, to fill out the combination, and
3) cull out duplicates and those violating certain facility constraints (e.g. no more than two new wells could be drilled from the northern platform).

The 150 examples set aside for testing the ANNs generalization performance were generated in the same fashion, except that sampling proceeded until exactly 50 2-, 3-, and 4-well combinations were obtained. This balancing by size is intended to avoid inadvertently biasing the test set in favor of any particular size, which can occur when random methods are applied to a relatively small sample.

By the standard formula for combinations of n elements taken r at a time, the total possible combinations of 1-, 2-, 3- and 4-well combinations are 25, 300, 2300, and 12,650, respectively. The entire knowledge base, including both training and testing examples, contained 25, 158, 184 and 182 combinations, respectively. The rate of inclusion of each of the 25 locations ranged from 0.11 to 0.14.

Carry out the Simulations - A key feature of the collection of examples generated in the sampling step is that they are independent of each other. The input to example B is not

dependent on the outcomes of example A. Consequently, they can be farmed out, either manually or in an automated fashion, to as many processors as the simulator's license allows. At the time the knowledge base simulations were conducted for this project, only two single-user licenses were available. So, the simulation step required about six weeks to complete. Given more licenses, this task could have been completed more quickly.

In contrast to typical reservoir modeling studies where detailed attention is paid to setting the simulation parameters and the analysis of outcomes on a case-by-case basis, both the creation of input files and the analysis of output is automated. General rules for assigning skin factors to injection locations, determining the layers in which a well would be completed, and setting facility constraints, together with appropriate simulation parameters, were obtained from members of the Pompano asset team who had been closely involved in the development and use of the numeric model. Given these rules and the list of well combinations to simulate, Perl scripts tailored input files for each run, launched the simulation, and extracted and saved information from each run's output.

Train ANNs to Predict Reservoir Performance. The architecture used for all ANNs in the Pompano project was a feedforward network, trained by the familiar backpropagation learning algorithm [9]. In this paradigm, a network is initialized with small random weights, as is illustrated in Fig 5. Training consists of presenting example inputs to the network and calculating the corresponding outputs, given the current values of the connection weights. The calculated output values are compared to the target values from the examples; and the connection weights are updated according to any of several learning algorithms to minimize the difference between calculated and target values on the next iteration. Over time, the connection weights associated with important relationships grow large and those associated with trivial relationships decay to zero. In the particular implementation used for the Pompano project, a conjugate gradient optimization method [24], employing the Polak-Ribiere weight update rule, was used to speed convergence and reduce the likelihood of becoming trapped in local minima. A sigmoid was used as the transfer function. To avoid overfitting of the network weights to idiosyncratic features of the training examples, batch updating of weights and a relatively short number, 300, of training epochs was employed.

The goal of training is to construct a network with maximal capacity to accurately generalize its predictions to previously unseen combinations. Accuracy is defined here as the square of the Pearson product-moment correlation, r^2, between the ANN's and the simulator's predictions for a given attribute on some set of examples. Training accuracy, then, is the r^2 between the ANN and simulator predictions on the examples in the training set. Testing or generalization accuracy is this same measure on the examples in the test set. Factors that are known to contribute to generalization include the complexity of the network as reflected in the number of connection weights, the size and composition of the training set, and the degree of noise in the

training/testing sets [25]. In the current study, noise in the usual sense of the term is not at issue since the examples are generated by mathematics, not nature. This is probably the main reason why all ANNs in the Pompano problem achieved very high levels of *training set* accuracy (r^2 >= .95), a necessary but not sufficient condition for generalization accuracy.

The issue of training set size, on the other hand, is much more problematic. The allocation of 400 combinations to the training set and 150 to the testing set, as described earlier, was based mainly on experience gained in two prior optimization studies conducted on a groundwater remediation problem. Although these numbers were thought to be low, given the greater degree of nonlinearity in 3-D multiphase flow, it was deemed to be preferrable to proceed with a manageable number of simulations and leave the question of the relationship between training/testing set sizes and predictive accuracy to later research efforts.

The third factor, network complexity, is addressed by the manner in which variations on a given network are constructed and tested. As illustrated by the simplified network in Fig. 5, the size of the input and output layers are fixed at 25 nodes and one node, respectively, these dimensions having been established as the minimum necessary to adequately represent the Pompano problem. Earlier efforts to express well locations in x-y coordinates to permit a network to make spatial interpolations produced greatly degraded predictive accuracy. So, the convention of employing a set of preselected locations that constitutes the domain about which questions can be asked has been followed in this work. To keep the architecture similarly stream-lined, networks are constructed to predict only one attribute at a time: 7-year cumulative oil production, 7-year cumulative gas production, and peak injection volume. The results of searches optimizing SNP over three years, which required 3-year versions of cumulative oil and gas production, proved to be uninteresting, because there was little performance spread between well combinations over such a short period of time. Consequently, the 3-year ANNs will not be discussed, except to illustrate an occasional point about neural network training and testing. One such point is that, since the knowledge base contained yearly performance data, it could be used to train ANNs over any desired time-frame up to the maximum of seven years.

The only variable architectural element, then, is the number of nodes in the hidden layer. This value of this attribute which best promotes generalization is determined empirically by training variant networks with anywhere from 1 to 10 hidden nodes and selecting the variant with the best test set (i.e. generalization) accuracy. The protocol for selecting the best possible ANN for a given predictive task cannot end there, however. Backpropagation training is, itself, a nonlinear optimization problem and suffers from vulnerability to entrapment in local minima in the error-surface, depending on the randomly-assigned initial values of the connection weights. The variance caused by those initial values is partly a function of the complexity of the input-output relationships being mapped and can also be reduced by increasing the size of the training

set. However, with the relatively small training/testing set sizes in the Pompano problem, some other procedure had to be developed to confront the initial-weights issue.

Fig. 6 illustrates the kind of initial-weights analysis that was performed. The graphs show mean test set accuracy, ± one standard deviation, over 25 different weight initializations for each hidden layer size from 1 to 10. The complete training of variant networks for each attribute (e.g. 3-year cumulative gas) required 250 training/testing cycles. The task was performed by a batch process that required a total of about one hour to complete, per attribute. The purpose of this exercise was to select a size for the hidden layer with not only the highest mean but also the smallest standard deviation, in an effort to identify the network architecture with the *best and most stable* generalization. Having narrowed the number of variants being considered to 25 by selecting the size of the hidden layer, the network chosen to participate in the searches was simply the variant with the highest test set accuracy.

Fig. 6 also shows how different the various attributes being predicted can be from each other. The easiest attribute to accurately predict is 7-year cumulative gas, as shown by its very high means and tiny standard deviations. Defying the usual rule-of-thumb that predictive accuracy declines with increasing time, 3-year cumulative gas shows slightly lower accuracies. The situation reverts to expectations with cumulative oil, however, 3-year performance being considerably easier to predict accurately than 7-year performance. These results underscore the critical point that the mapping of inputs to outputs by the ANNs is an empirical procedure. The complexities and outcomes of the mapping is a function of the particular examples in the training and test sets and do not necessarily reflect more general physical principles.

Search for Optimal Well Combinations. Although the trained nets can be coupled with a variety of search techniques, the genetic algorithm (GA) and simulated annealing (SA) methods were selected for their robustness and flexibility. Like all optimization drivers, these techniques are highly sensitive to some of the parameters guiding their search and relatively insensitive to others. The parameter settings used in the Pompano project have been determined by extensive trial-and-error experimentation. A discussion of these methods is given below.

Genetic Algorithm - Given the attention that has been paid to GA applications in recent years, readers are probably familiar with the basic mechanisms of and rationale for this family of search techniques. Consequently, this section will mainly address the specific procedures chosen for implementation in the current study. Excellent introductions can be found in Goldberg [10] and Michalewicz [26]. Goldberg is the source for all information concerning the GA presented below, unless otherwise noted.

The upper portion of Table 1 presents a summary of parameters and procedures used in the current GA. The 25 well locations which form the decision variables are represented in the GA as a string of 25 bits, each of which can either be on or off. Both the spatial location and flow rate of each well is fixed and implicit in the representation. The order of the well locations in the string is indicated by their identification numbers in Fig. 3. That numbering is arbitrary, as is their location in the bit-string.

The search is initialized with a set of 100 well combinations. In fact, this initial population is simply a random subset of the 150 cases in the ANN testing set. The population size of 100 chosen for the current study is a fairly small value. In water resources applications, values have ranged from 64 [27] to 300 [28] and even up to 1,000 [29]. A larger population helps maintain greater diversity but does so at considerable computational cost when the full model is being used to generate performance predictions.

The basic cycle of the GA is as follows. The initial population of 100 well combinations is evaluated according to an objective function, SNP in this case. A new generation of 100 combinations is created from the old population by means of three mechanisms: selection, reproduction, and mutation. The new population is then evaluated according to the objective function; and the entire process is repeated until some termination criterion is reached. The manner in which the three mechanisms have been implemented is as follows:

1) Selection - This mechanism determines which members of the current generation will be selected for carry-over, in one form or another, to the new generation. To make sure that the highest-ranking combinations are not lost to the population through accidents of selection and crossover, the top three combinations are copied over to the new generation intact. The remaining 97 slots in the new population are filled by a form of sexual reproduction, a process for which parents must be selected.

The most popular method of selection is the roulette wheel, in which each member's likelihood of being selected for reproduction is the ratio of its own performance score to the total performance score of the population. The larger a given member's score is in relation to the other members', the larger portion of the roulette wheel it occupies, increasing the odds that the member will be selected one or more times for reproduction. When large discrepancies exist in the scores of individual members, the members with the higher scores come to dominate the population too quickly. Conversely, when differences between members become very small, the selection process becomes random. To avoid these cases, the current GA employs selection based on the combinations' rank order [30] rather than their proportional scores. Combinations are selected by sampling from a uniform distribution over the ranks, with a bias factor of 1.5 serving to favor high-ranking combinations over lower-ranked combinations.

Selections for reproduction are made, two at a time, to obtain parent combinations from which a child combination will be formed. This process is repeated until 97 children have been generated. The same combination may constitute both members of the pair, in which case the child is simply a clone of the parent.

2) Reproduction (Crossover) - The most common form of reproduction is single-point crossover. Child combinations are constructed by breaking the parent combinations apart at some randomly selected crossover position in the bit-string and joining segments from each parent. For example, given two parents in a 5-bit problem (0 1 0 0 0 and *1 1 0 1 1*) and a crossover point of 2, two different children could be constructed (0 1 *0 1 1* and *1 1* 0 0 0).

Creating new combinations from "chunks" of old ones makes the most sense when proximity in the bit-string is important. That is, the proximity of wells in the bit-string should reflect one or more dimensions of relatedness in the physical problem it represents. This is not necessarily the case in the Pompano problem. In fact, the earlier groundwater studies employing the GA had discovered a "sticky" well problem. That is, particular wells kept appearing in the optimal solutions sets whose individual contributions to the efficiency of remediation were minimal but which were adjacent in the bit-string to wells making major contributions. To break up these spurious associations, a different reproductive mechanism, uniform crossover, is used [31]. In this method, the value of each bit in the child string is set independently of every other bit. A coin-toss at each bit-position determines from which parent the child will inherit the value for that particular bit. The exchange probability can be biased to favor the fitter parent, if any; but in this study the exchange probability is kept at an impartial 0.5.

3) Mutation - Mutation is a way to maintain diversity in a population by arbitrarily changing the values of bits in the child combinations according to some rate, often the inverse of the population size. A high mutation rate can undermine the effects of crossover; a low one limits the introduction of "novelty" into the population. For this study, the inverse rule yields a mutation rate of 0.001.

Simulated Annealing - Like the GA, SA techniques are based on an analogy to a natural process. Instead of Darwinian concepts of evolution, which are the foundation of the GA, SA is based on an analogy to the cooling of materials in a heat bath. The fundamental idea is that if the amount of energy in a system is reduced very slowly, the system will come to rest in a more perfect state than if the energy is reduced quickly. When translated into terms pertinent to optimization, the energy in the system refers to the tolerance for pursuing apparently poorer solutions in an effort to avoid being trapped in local minima. As the search proceeds, this tolerance is slowly reduced until the search converges to a final optimal solution. SA algorithms

have appeared in several applications [32-35]. A highly readable introduction to the subject can be found in Dowsland [36], which is also the source for the material discussed below, unless otherwise noted.

The SA parameters employed in the current study are given in the lower portion of Table 1. SA represents a return to single-point search, in contrast to the multiple-point or population-based search of the GA. At every step, there is only one new well combination being compared to the current combination. The initial combination represents the starting point for search. In this implementation, the initial combination is the no-injection case. Trial and error experimentation with the algorithm has shown that the initial starting point has only a small effect on the duration of search. The current study's annealing algorithm, adapted from the standard algorithm as presented in Dowsland [36], proceeds as follows:

Set the current combination c = initial combination

Set the current energy in the system t = initial temperature

Select a temperature decrement function α

Repeat

  Repeat

    Generate a new combination n in the neighborhood of c

    δ = fitness(n) - fitness(c)

    if $\delta > 0$ then $c = n$

    else

      generate a random value x uniformly in the range (0,1)

      if $x <$ exp(-δ/t) then $c = n$

  Until the iteration counter = iterations/temperature

  Set $t = \alpha(t)$

Until termination criteria are met

The purpose of the temperature parameter in the algorithm is to control the tolerance for accepting a newly generated combination n as the current combination c, even when its performance score is lower than the current combination's score. If the new combination's score is greater than the current combination's, it is always accepted as the new current combination. If not, there is a probability of accepting it anyway that is a function of the current temperature t in the system, leavened by the magnitude of the difference δ between the two scores. The initial

temperature and the range over which it is allowed to vary are empirically determined parameters. The experimenter decides, in advance, what overall percentages of poorer combinations it is desirable to accept in the initial and final stages of search and adjusts the temperature range until those percentages are achieved.

On the other hand, the temperature decrement or cooling function and the number of iterations per temperature are parameters that have received more attention in the literature. As was mentioned earlier, the rate of cooling has considerable impact on the likelihood of converging to an optimal solution. The function chosen for the current study, a geometric decrement function with a decrement factor of 0.9, is one of the two most widely used approaches. The issue of how many iterations to perform at a given temperature level has been the subject of considerable analysis in certain applications [32, 34]. While theory suggests that extremely large values for this parameter should be used to guarantee that the algorithm is given an adequate opportunity to sample the search space, experimentation with this parameter indicates that much smaller values, 10-100 times the number of decision variables, can be employed. At a minimum, this rule of thumb would imply that iterations/temperature should be set to 250 for the current study. Instead, a very small value, 50, has been selected, mainly to permit more timely comparisons between ANN- and simulator-based searches.

The algorithm listed above glosses over an important function, that of generating the new combination from the nieghborhood of the current combination. This is another domain-dependent decision because the manner in which valid new combinations can be constructed from old ones is a function of the problem representation. In the current implemention, the temperature parameter is used here, too, to control the extent to which the new combination can vary from the current combination. This is equivalent to controlling the size of the local nieghborhood being searched at a specified temperature level. Initially, the number of well locations in the current combination that will be switched is determined by randomly selecting an integer from 1 to 10. The particular locations to alter are then selected at random from the available 25 locations, subject to the usual facility constraints described earlier, until the prespecified number of locations in the current combination have had their status changed from *on* to *off* or vice versa. As temperature decreases, the maximum number of locations that can be potentially changed is reduced from 10 to 1.

A small departure, also not shown above, from the serial nature of the standard algorithm has been implemented. According to the standard algorithm, the current combination at the end of processing at a given temperature level is not necessarily the highest-scoring combination encountered during the 50 iterations at that level, because that there is a certain probability that an inferior new combination will replace the current combination. However, the algorithm implemented in the current study remembers the best combination ever encountered and makes it the current combination before proceeding to the next temperature level. This is somewhat akin

to the practice in the GA of preserving the top combinations from one generation to the next so that they are not lost through the vissicitudes of selection and crossover.

Procedures Common to Both GA and SA Searches - Termination criteria in optimization are usually based on some notion of convergence to a single best solution. In keeping with the philosophy of heuristic search, however, the current study is more interested in generating *sets* of near-optimal solutions rather than a single best solution. This goal is achieved by tying termination criteria to the performance score of the population, in the case of the GA, or the temperature, in the case of SA, rather than the performance of the highest-ranking individual combination. Search terminates when the mean population or temperature performance score fails to increase over five consecutive generations/temperatures or some maximum number of generations/temperatures have elapsed, whichever comes first. The maximum number of the GA generations is 25; the maximum number of SA temperatures was reduced to 16, to prevent over-long searches when the ANN-SA vs. VIP®-SA comparison was being conducted. At the end of every generation/temperature, combinations with scores above a predetermined cut-off are saved to a file. The top-ranked unique combinations in this file become the set of near-optimal solutions.

The outcome of search in both the GA and SA is influenced by the particular randomly-based choices that are made. To avoid the possibly idiosyncratic results of any single search, the results of each search in the current study (with one exception, which will be noted below) actually consist of combined results from 10 searches, each with a different seed initializing the pseudo-random number generator.

Verify Optimal Combinations with the Simulator. In an actual engineering application of the ANN-GA/SA methodology, the asset team may choose to only submit a handful of well combinations to the simulator. For this demonstration project, however, the top 25 well combinations from the near-optimal set were submitted for verification. The resulting simulator predictions of 7-year oil and gas production and peak injection volume are used to recalculate the SNP. The updated SNPs are intended to become the measure for subsequent analysis and decision-making.

Search Results

The results of various efforts to identify optimal well combinations to maximize SNP over seven years are shown in Tables 2-6. Throughout, the production figures for the no-injection baseline case serve as the standard against which alternative scenarios are judged. All values appearing in the tables are reported in increments/decrements of the appropriate unit (e.g. dollars, mmcf). SNP is calculated according to the cost estimates and definitions described earlier. Scenarios are

designated by a list of the identification numbers (see Fig. 3) of the wells making up the combination.

Context Scenarios. It is useful to begin by considering the performance of some simple conversions of existing producers in the northern section to injection locations. There is considerable appeal to pursuing such scenarios, in part because knowledge of the reservoir is much greater in the vicinity of an existing producer than around the new injection locations. The effects of converting the four northern producers which survived the initial screening of locations (according to the criterion that their effect, singly, on total hydrocarbon production over seven years must exceed the no-injection baseline) are shown in Table 2. When the more complex SNP performance measure is used, two of the wells now show a negative impact on total field productivity and the postive influence of the other two is minimal. The best conversion, well 12, produces only a 1.72% improvement over baseline performance. This result illustrates the sensitivity of outcomes to the particular performance measure being used and suggests that many different measures should be used to evaluate scenarios for field development.

Best in Knowledge Base. The next most obvious tactic is to query the 550-case knowledge base to identify the well combinations which yield the highest SNPs. The attraction of this tactic is that the oil production, gas production, and peak injection volume inputs to the SNP calculations come directly from the simulator, without any estimation errors introduced by the ANNs. The drawback is that results are limited to well combinations already in the knowledge base. As Table 4 shows, the information in the knowledge base alone makes a considerable improvement in expected performance of the field over the simple single-well conversions of Table 2. The best combination, consisting of wells 7 and 9, shows an 11.11% improvement over the baseline SNP.

ANN-GA Search Results. The reason for going to the extra effort of implementing an actual search for optimal well combinations is that there may be combinations not sampled in the knowledge base which have superior performance characteristics. A directed search technique can usually identify peak performers which a random sampling may miss. Since the time required to train ANNs and conduct the searches is small (at least once the methodologies are mastered) relative to the time required to create the knowledge base, there is ample reason to proceed.

The entire 10-cycle GA search required less than an hour on the same class of workstation used to perform the simulations. All well combinations with estimated SNP's above a certain cut-off were saved and combined for post-processing. The top 25 well combinations from this pool were submitted to the simulator to verify the oil, gas, and peak injection numbers and calculated an updated SNP. The 16 combinations whose updated SNP exceeded the best in the

knowledge base are shown in Table 4. The best combination, 7-16-24, yields a 13.5% improvement over baseline.

In addition to fulfilling the final step in the ANN-GA/SA process, that of verifying the optimal set of well combinations so that engineering decisions can be made on the best-available information, the data generated by the verification runs provide an opportunity to assess the final-stage accuracy of the ANNs. The generalization accuracies, expressed as the squared correlation between ANN and simulator predictions, of the 7-year oil, 7-year gas, and peak injection volume ANNs on the 150-case test set were .81, .98, and .99, respectively. To the extent that the test set is a good, if low-resolution, representation of the total space over which the search might roam, these numbers indicate excellent generalization for the gas and peak injection ANNs and borderline-acceptable generalization for the oil ANN. This does not mean, however, that the ANNs' level of accuracy will be maintained during the final stages when small subregions of the search space are being searched at high resolution. It is to be expected that ANNs trained and tested on a coarse sampling will lose accuracy when required to make fine-grained distinctions. And, in fact, the squared correlations between ANN and simulator predictions on the top 25 well combinations generated by the ANN-GA search on 7-year oil, 7-year gas, and peak injection volume were .39, .38, and .96, respectively. Furthermore, the correlation between the SNP estimates based on the ANN predictions and the updated SNP figures based on the simulator-verified numbers were virtually zero. What this suggests is that the GA, using SNP calculations based on the relatively coarse-grained ANN predictions, is able to locate appropriate regions where optimal combinations lie and to identify several near-optimal candidates; but only the simulator itself can sort out the relative ranking among that final set of candidates.

ANN-SA Search Results. One issue that needs to be explored is the extent to which the results in Table 4 are a function of the ANNs or a function of the GA. To address this question, an ANN-SA search was conducted, holding all procedures used in the previous search constant except for the substitution of the SA search method. The results, shown in Table 5, are a very clear indication that the final set of well combinations is being determined by the ANNs. The shift to a different search method made almost no difference to the final set of well combinations.

VIP®-SA Search Results. One last question, given that the ANNs' final stage accuracy does deteriorate, is whether superior well combinations are being missed because the ANNs, rather than the simulator, are supplying the predictions which are influencing the direction of search. The only way to answer this question is to conduct a search in which the simulator is called each time a new well combination is being evaluated by the search engine. A VIP®-SA search, the results of which are shown in Table 6, was conducted in the following manner: Because the simulator was called to supply the oil, gas, and peak injection data needed for the SNP calculations and each call required an average of 3.5 hours, only three rather than 10 repetitions of

the search were performed. Three workstations ran in parallel, sharing a common cache of results so that no time would be wasted on duplicate calls for the same well combination. The SA engine was chosen over the GA because the former method tends to converge more quickly. Even so, the three searches required several weeks to complete on each workstation, involved 936 unique calls to the simulator (or 3276 total computational hours) and matched but did not beat the best well combination (7-16-24) found by the ANN-SA search. In other words, using the simulator directly in the search did not improve the quality of results for this particular problem; it merely took an inordinate amount of computational time. Furthermore, results saved from these runs are not likely to be re-usable in new searches. The well combinations which dominate for the current definition of SNP will not necessarily appear at all if another set of cost estimates is used.

Summary and Conclusions

Scope of Results. The purpose of the Pompano project has been to apply a methodology originally developed to optimize the placement of wells in groundwater remediation to a problem in reservoir management. Given this demonstration focus, the project has sought to illustrate the improvements in decision-making which can be achieved with only minimal adaptation of methods from the earlier work. The following conclusions seem to be warranted by the results shown in Tables 2-6:

Current practice in the industry is to treat the reservoir simulator as a tool for detailed analysis of the reservoir. Members of the asset team propose a small number of scenarios for well-field development based on the available reservoir characterization data, any existing production data, and their own knowledge and experience. These scenarios are submitted to the simulator, with results confirming or refuting the team's proposals and possibly suggesting new design variations to explore. Because the emphasis is on detailed examination of results, the total number of scenarios that are likely to be considered in this approach is on the order of "tens".

Using the archive of simulations as a database, the well-field optimization project has introduced a change of perspective, expanding the scope of study from "tens" to "hundreds". The simulator is now viewed as a tool for providing rapid answers to a variety of engineering and management questions. Querying the database of simulations has highlighted the considerable increase in performance that may possibly be achieved by switching from an approach to injection based on converting one or more existing producers to one involving the drilling of three to four new injectors, despite the increased capital and operating expenses associated with this latter approach.

Even greater value is mined from the reservoir simulator when the archive of simulations is used in its second capacity: as a source of examples for training and testing ANNs. We have trained ANNs to predict peak injection volumes and volumes of produced oil and gas over seven

years of injection. The rapid estimates of these quantities provided by the ANNs are fed into simple net profit calculations, which in turn are used by the GA or SA to evaluate the effectiveness of different well-field scenarios. The search engine explores scenarios not contained in the original archive of simulations, expanding the scope of study from hundreds of scenarios into the "thousands". This expansion has enabled the identification of new scenarios which exceed the simple net profits of the best scenarios found by simply querying the database of simulations

Issues. Both substantive and methodological issues have been raised in the course of the pilot project's activities:

Substantive Interpretation of Results - In the discussion of results in Tables 2-6, emphasis was placed on the best performing scenario located by each method. However, the results actually consist of *sets* of near-optimal scenarios which can be analyzed in an effort to better understand the underlying physical reasons why these scenarios are optimal answers to a particular management question. For example, an examination of the top 25 well combinations from the ANN-GA search found that well 7 figured in 100% of the combinations, followed at a distance by well 24 (32%), well 11 (28%), and well 16 (24%). One might speculate that well 7 has a larger sweep of neighboring producers that are important to production over the seven year time-frame. The other popular wells may be reflective of more conventional wisdom regarding the desirability of raising pressures near the boundaries of the reservoir. Given results from several searches addressing different management questions (e.g. varying economic parameters and time-frames or narrowing the focus to wells of special interest), the asset team has the opportunity to build a body of operating principles about the field, some of which may transfer to other fields, as well.

ANN Accuracy Issues - An ANN's generalization accuracy has to be examined in two ways. First, there is the question of how accurately it makes predictions over the entire space in which predictions might be called for by the search engine. In the Pompano problem, these initial accuracies were estimated by correlating ANN and simulator predictions on the 150 cases in the test set. Accuracy varied considerably by both the attribute to be predicted (oil vs. gas, for example) and the time-frame over which the prediction was being made (three vs. seven years). It appears, however, that despite errors of estimation, the search engines still gravitated to the specific regions where optimal well combinations were to be found and generated several near-optimal candidates. It was then the job of the simulator to sort out the proper rankings among the "finalists". One possible drawback to this technique, that ANN-introduced errors would cause the search to completely miss the best combinations, proved not to be true for this particular problem. In a comparison search where the simulator itself was called upon to provide predictions as demanded by the search engine, the best combination located in the ANN-based searches was matched but not beaten. Still, this outcome might not be born out on other

management questions or over longer time-frames. Consequently, a critical research task is to identify strategies for improving both initial and final accuracy. An example of a technical strategy which might improve accuracy involves simplifying the prediction task for the ANNs by training separate nets for the 2-, 3-, and 4-well combinations. A more knowledge-based approach to the problem would involve supplying more information about the reservoir to the nets, in the form of more inputs. For example, instead of describing a well combination as a set of binary inputs, the local average permeability of each well which will be turned on in the combination could serve as the inputs.

Uncertainties in the Underlying Model - A third concern that has been raised in the technical reviews of this project involves uncertainties associated with the underlying reservoir simulator. So far, a single model of the reservoir has been taken as a kind of norm or best-bet on which to base reservoir management decisions. However, this approach tends to gloss over the possibility that reasonable alternatives to the normative model exist which may greatly affect the optimal solutions to management questions. Decision-makers are better served if they are presented with at least some indication of how great a variation is introduced by considering these alternatives.

The problem of estimating and managing model uncertainties is huge and will not be solved anytime in the near future. There are, however, incremental strategies for incorporating aspects of uncertainty analysis into the ANN-GA/SA methodology at different stages of the optimization process. A very simple strategy might be to rank each well location by the relative certainty of the physical properties in its vicinity. The objective function being optimized would contain a penalty term based on that rank, which will reflect the informational-risk associated with including that well in the scenario. A much more laborious approach would be to create separate knowledge bases for a small set (e.g. three) of geologically-reasonable alternative models and carry out the entire process separately for each one, comparing results for common locations.

Acknowledgements

This work has been funded by DeepLook, a consortium of major oil and service companies, dedicated to collaborative research in the area of locating and exploiting bypassed hydrocarbons. Special thanks are due to Landmark Graphics, Inc., for use of their reservoir simulator, VIP®. This work has been performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-408.

References

[1] L. L. Rogers and F. U. Dowla, "Optimization of groundwater remediation using artificial neural networks and parallel solute transport modeling," *Water Resources Research*, 1994, 30(2), 457-481.

[2] L. L. Rogers, F. U. Dowla, and V. M. Johnson, "Optimal field-scale groundwater remediation using neural networks and the genetic algorithm," *Environmental Science & Technology*, 1995, 29(5), 1145-1155.

[3] B. J. Wagner, "Recent advances in simulation-optimization groundwater management modeling," *Reviews of Geophysics, Supplement*, U. S. National Report to International Union of Geodesy and Geophysics, 1991-1994, 1995, 1021-1028.

[4] G. P. Karatzas and G. F. Pinder, "Groundwater management using numerical simulation and the outer approximation method for global optimization," *Water Resources Research*, 1993, 29(10), 3371-3378.

[5] B. J. Wagner and S. M. Gorelick, "Reliable aquifer remediation in the presence of spatially variable hydraulic conductivity: From data to design," *Water Resources Research*, 1989, 25(10), 2211-2226.

[6] N. Chan, "Robustness of the multiple realization method for stochastic hydraulic aquifer management," *Water Resources Research*, 1993, 29(9), 3159-3167.

[7] D. E. Dougherty, "Hydrologic applications of the connection machine CM-2," *Water Resources Research*, 1991, 27(12), 3137-3147.

[8] A. F. B. Tompson, S. F. Ashby, R. D. Falgout, S. G. Smith, T. W. Fogwell, and G. A. Loosmore, "Use of high performance computing to examine the effectiveness of aquifer remediation," In *Computational Methods in Water Resources X, Volume 2*, A. Peters, ed., Boston, MA: Kluwer Academic Publishers, 1994, pp. 875-882.

[9] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning internal representations by error propagation," In *Parallel distributed processing: explorations in the microstructure of cognition, Vol. 1: Foundations*, D. E. Rumelhart and J. L. McClelland, eds., Cambridge, MA: MIT Press, 1986, 318-362.

[10] D. E. Goldberg, *Genetic algorithms in search, optimization, and machine learning*, Reading, MA: Addison-Wesley, 1989.

[11] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing", *Science,* 1983, 220, 671-680.

[12] C. R. Reeves, "Introduction." In *Modern heuristic techniques for combinatorial problems*, C. R. Reeves, ed., Halstead Press, New York, N.Y., 1993, 1-19.

[13] E. A. Breitenbach, "Reservoir simulation: state of the art", *JPT,* September, 1991, 1033-1036.

[14] P. Behrenbruch, "Offshore oilfield development planning", *JPT,* August, 1993, 735-743.

[15] A. Aadland, O. Dymes, S. R. Olsen, and O. M. Dronen, "Statfjord field: field and reservoir management perspectives", *SPE Reservoir Engineering,* August, 1994, 157-161.

[16] S. D. York, C. P. Peng, and T. H. Joslin, "Reservoir management of Valhall field, Norway", *JPT,* August, 1992, 918-923.

[17] J. Kikani and T. D. Smith, "Recovery optimization by modeling depletion and fault block differential pressures at Green Canyon 110", SPE 36693, in *Proc. SPE Annual Technical Conference and Exhibition: Reservoir Engineering,* October 6-9, 1996, Denver, CO. Tulsa, OK: SPE, 1996, 157-170.

[18] G. Coskuner and B. Lutes, "Optimizing field development through infill drilling coupled with surface network: a case study of low permeability gas reservoir", SPE 36705, in *Proc. SPE Annual Technical Conference and Exhibition: Reservoir Engineering,* October 6-9, 1996, Denver, CO. Tulsa, OK: SPE, 1996, 273-284.

[19] M. Kumar and V. M. Ziegler, "Injection schedules and production strategies for optimizing steamflood performance", *SPE Reservoir Engineering,* May, 1993, 101-108.

[20] S. I. Aanonsen, A. L. Eide, L. Holden, and J. O. Aasen, "Optimizing reservoir performance under uncertainty with application to well location", SPE 30710, in *Proc. SPE Annual Technical Conference and Exhibition: Reservoir Engineering,* October 22-25, 1995, Dallas, TX. Tulsa, OK: SPE, 1995, 67-76.

[21] R. K. Wackowski, C. E. Stevens, L. O. Masoner, V. Attanucci, J. L. Larson, and K. S. Aslesen, "Applying rigorous decision analysis methodology to optimization of a tertiary recovery project: Rangely Weber sand unit, Colorado", SPE 24234, paper presented at the SPE Oil & Gas Economics, Finance and Management Conference, April 28-29, 1992, London, U. K.

[22] H. Fujii and R. N. Horne, "Multivariate optimization of networked production systems", SPE 27617, paper presented at the European Production Operations Conference and Exhibition, March 15-17, 1994, Aberdeen, U. K.

[23] A. C. Bittencourt and R. N. Horne, "Reservoir development and design optimization", in *Proc. SPE Annual Technical Conference and Exhibition: Reservoir Engineering*, October 1997, San Antonio, TX. Tulsa, OK: SPE, 1997, 554-558.

[24] E. M. Johansson, F. U. Dowla, and D. M. Goodman, "Backpropagation learning for multi-layer feed-forward neural networks using the conjugate gradient method," *Int. J. of Neural Systems*, 1992, 2(4), 291-301.

[25] A. Atiya and C. Ji, "How initial conditions affect generalization performance in large networks." *IEEE Trans. on Neural Networks*, 1997, 8(2), 448-451.

[26] Z. Michalewicz, *Genetic algorithms + data structures = evolution programs*, 2nd ed., Springer-Verlag, New York, N. Y., 1992.

[27] D. C. McKinney and M. D. Lin, "Genetic algorithm solution of groundwater management models." *Water Resources Research*, 1994, 30(6), 1897-1906.

[28] B. J. Ritzel, J. W. Eheart, S. and Ranjithan, "Using genetic algorithms to solve a multiple objective groundwater pollution containment problem." *Water Resources Research*, 1994, 30(5), 1589-1603.

[29] S. E. Cieniawski, J. W. Eheart, S. and Ranjithan, "Using genetic algorithms to solve a multiobjective groundwater monitoring problem." *Water Resources Research*, 1995, 31(2), 399-409.

[30] D. Whitley, D, "The GENITOR algorithm and selection pressure: Why rank-based allocation of reproductive trials is best." In *Proc., 3rd Int. Conf. on Genetic Algorithms*, J. D. Schaffer,ed., Morgan Kaufman, San Mateo, Calif., 1989, 116-123.

[31] G. Syswerda, "Uniform crossover in genetic algorithms." In *Proc., 3rd Int. Conf. on Genetic Algorithms*, J. D. Schaffer, ed., Morgan Kaufman, San Mateo, Calif., 1989, 2-9.

[32] D. E. Dougherty and R. A. Marryott, R. A, "Markov chain length effects on optimization in groundwater management by simulated annealing." In *Computational methods in geoscience*, W. E. Fitzgibbon and M. F. Wheeler, eds., Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1992, 53-65.

[33] G. Christakos and B. R. Killam, "Sampling design for classifying contaminant level using annealing search algorithms." *Water Resources Research*, 1993, 29(12), 4063-4076.

[34] R. A. Marryott, D. E. Dougherty, and R. L. Stollar, "Optimal groundwater management. II. Application of simulated annealing to a field-scale contamination site." *Water Resources Research*, 1993, 29(4), 847-860.

[35] D. M. Rizzo and D. E. Dougherty, "Design optimization for multiple managment period groundwater remediation." *Water Resources Research*, 1996, 32(8), 2549-2561.

[36] K. A. Dowsland, "Simulated annealing." In *Modern heuristic techniques for combinatorial problems*, C. R. Reeves, ed., Halstead Press, New York, N.Y., 1993, 20-69.

TABLE 1--SEARCH PARAMETERS

Genetic Algorithm (GA)

Population Size	100
Initial Population	Randomly generated patterns
String Length	25 (one bit per well)
Selection for Mating:	
Method	Rank order
Selection Bias Factor	1.5
Crossover:	
Method	Uniform
Exchange Probability	0.5
Mutation Rate	0.001
Termination Criterion	5 generations without exceeding peak mean fitness, 25 generations maximum

Simulated Annealing (SA)

Initial Pattern	All 25 wells off
Cooling Schedule:	
Iterations/Temperature	50
Decrement	0.9 (i.e. $t_{n+1} = t_n - 0.9t_n$)
Starting Temperature	0.25
Neighborhood Search:	
Perturbations at t_1	Up to 10 randomly selected wells
Perturbations at t_{25}	1 randomly selected well
Reduction Schedule	Same as cooling schedule
Termination Criterion	5 temperatures without exceeding peak mean fitness, 16 temperatures maximum

TABLE 2--PERFORMANCE MEASURES OF SELECTED SINGLE-WELL CONVERSION SCENARIOS RELATIVE TO THE NO-INJECTION CASE

Scenario	Simple Net Profit (millions)	Oil (MSTB)	Gas (MMCF)	Costs (millions)
No Injection	0.00	0	0	0.00
Well 12	-5.75	+3,227	-17,223	+12.72
Well 13	+14.68	+5,696	-24,353	+12.72
Well 14	+3.59	+3,859	-17,400	+12.72
Well 15	-10.18	+688	-3,250	+12.72

TABLE 3--TOP 10 WELL COMBINATIONS FROM THE 550-CASE KNOWLEDGE BASE RANKED BY THEIR IMPROVEMENT OVER BASELINE ON SIMPLE NET PROFIT

Scenario	Simple Net Profit (millions)	Oil (MSTB)	Gas (MMCF)	Costs (millions)
7-9	+95.08	+11,899	-11,139	+61.5
7-24	+94.95	+11,715	-11,057	+61.5
6-19-21-24	+91.62	+19,007	-12,362	+94.5
7-16-18-24	+91.20	+15,152	-18,863	+96.5
6-7-10-16	+87.67	+14,307	-15,832	+94.5
7-11-15-16	+86.31	+14,176	-17,675	+89.2
7-25	+86.10	+11,233	-10,600	+61.5
6-7-10	+85.44	+12,275	-10,728	+78.0
6-7-9	+81.22	+12,676	-14,103	+80.0
7-11-21-23	+78.90	+13,666	-14,567	+96.5

TABLE 4--VIP® Verified Simple Net Profit (SNP) of Well Combinations from the ANN-GA Search Which Exceed the SNP of the Best Combination in the Knowledge Base

Scenario	Simple Net Profit (millions)
7-16-24	+115.54
7-16-23	+114.85
7-11-16	+113.89
1-7-24	+109.76
7-19-24	+109.76
6-7-24	+108.53
6-7-23	+107.96
7-16-25	+107.88
7-11-19	+106.94
1-7-11	+105.66
6-7-11	+104.02
7-20-24	+101.05
7-11-20	+99.14
7-9-20	+97.66
6-7-25	+96.83
7-9-16	+95.42

TABLE 5--VIP® Verified Simple Net Profit (SNP) of Well Combinations from the ANN-SA Search Which Exceed the SNP of the Best Combination in the Knowledge Base

Scenario*	Simple Net Profit (millions)
7-16-24	+115.54
7-16-23	+114.85
7-11-16	+113.89
1-7-24	+109.76
7-19-24	+109.76
6-7-24	+108.53
7-16-25	+107.88
7-11-19	+106.94
1-7-11	+105.66
6-7-11	+104.02
6-11-19	102.68
7-20-24	+101.05
7-11-20	+99.14
7-9-20	+97.66
7-9-16	+95.42

*Bold face scenarios were also located by the ANN-GA search (see Table 4).

TABLE 6--Simple Net Profit (SNP) of Well Combinations from the VIP®-SA Search Which Exceed the SNP of the Best Combination in the Knowledge Base

Scenario*	Simple Net Profit (millions)
7-16-24	+115.54
1-7-24	+109.76
7-19-24	+109.76
1-7-23	+108.73
6-7-23	+107.96
7-16-25	+107.88
7-16-20-24	+103.73
7-20-24	+101.05
7-9-20	+97.66
1-7-23-24	+97.61
1-7-11-23	+97.28
1-7-20-24	+97.28
6-7-10-23	+97.02
6-7-25	+96.83
7-16-20-25	+95.83
1-7-10-24	+95.27

*Bold face scenarios were also located by the ANN-SA search (see Table 5).

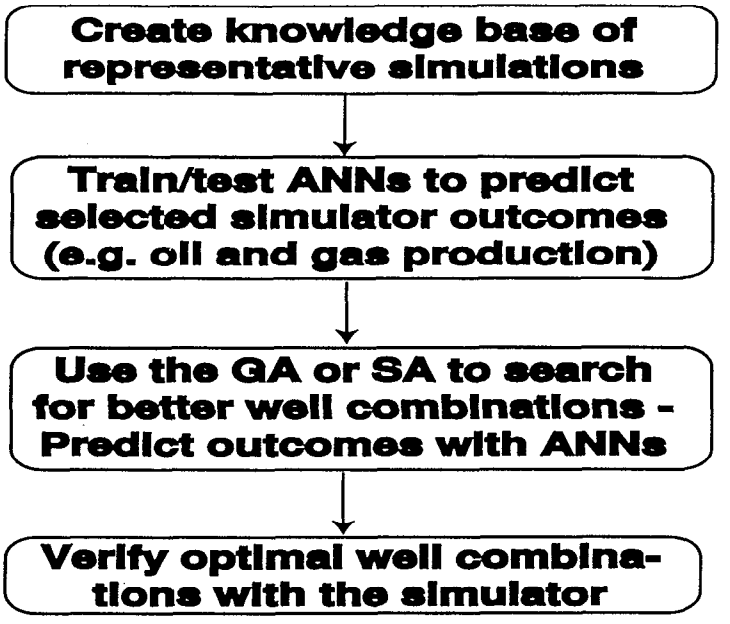

Fig. 1--Components of ANN-GA/SA Methodology

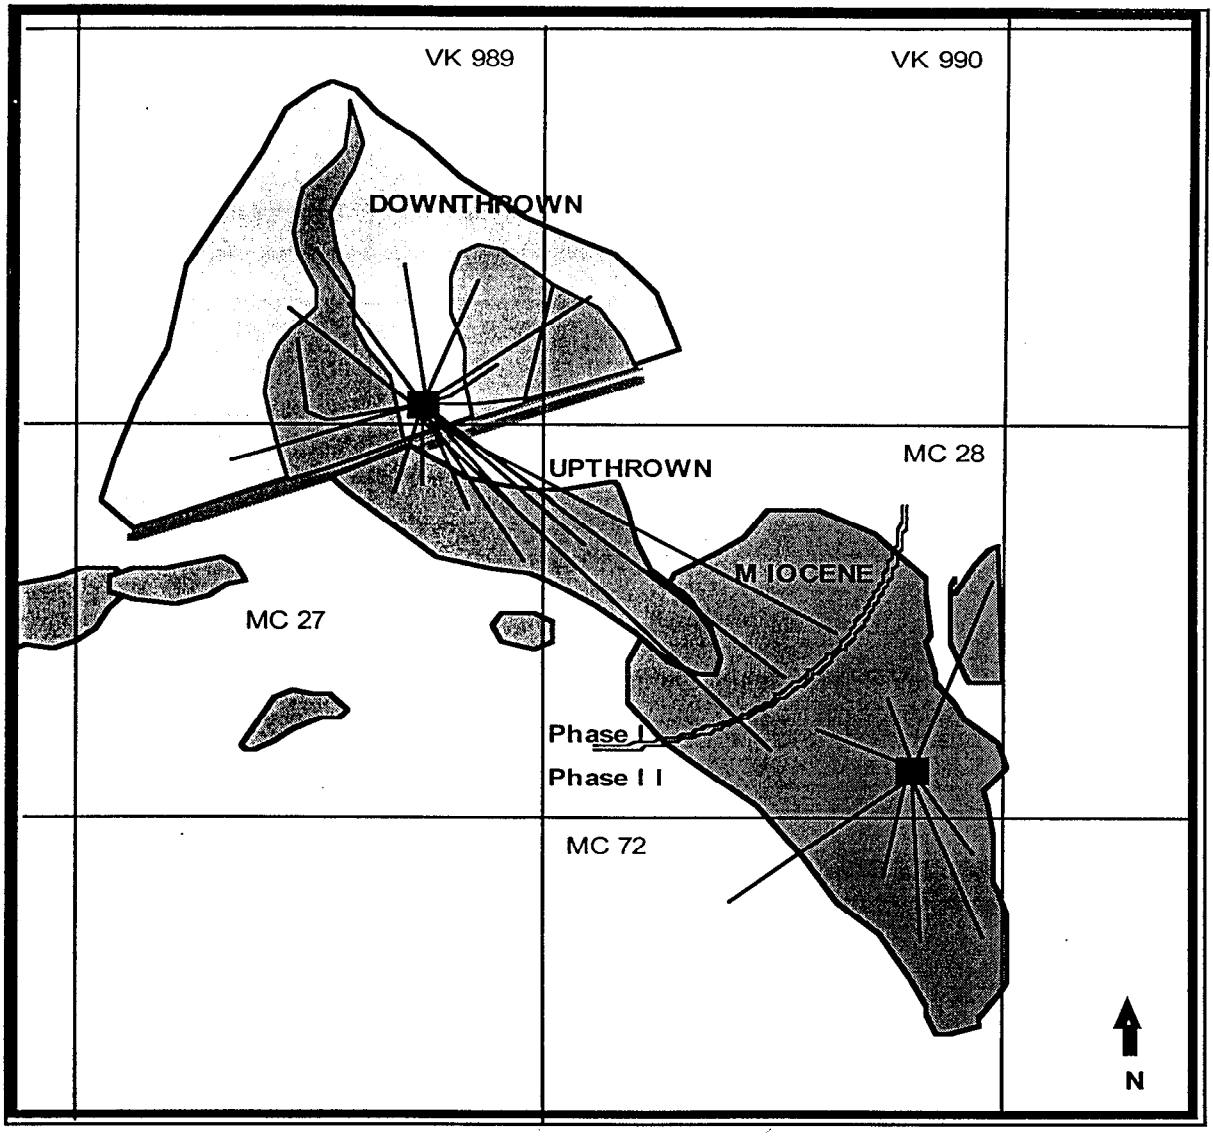

Fig. 2--Pompano Field

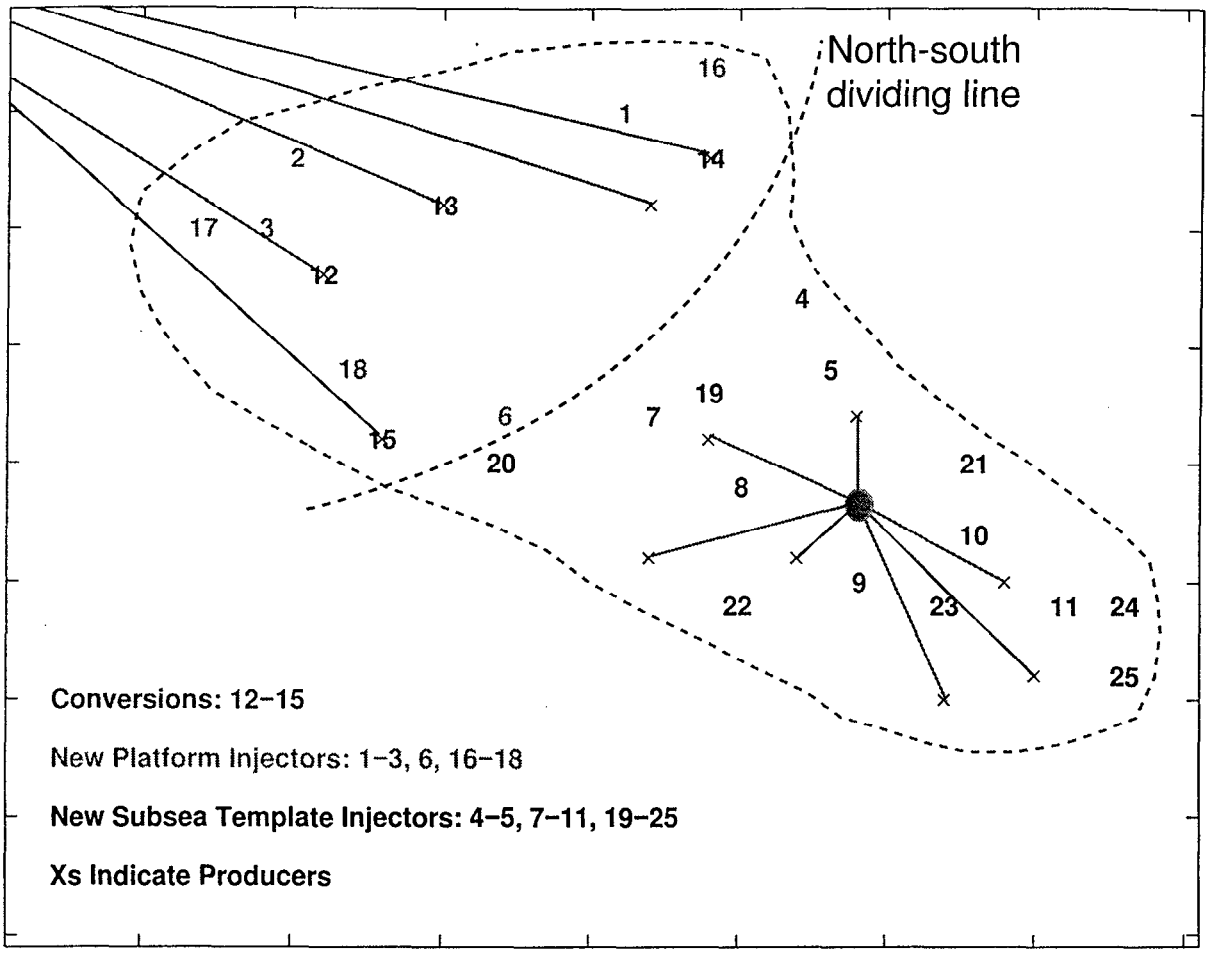

Fig. 3--Production and Injection Well Locations in the Northern and Southern Portions of the Miocene

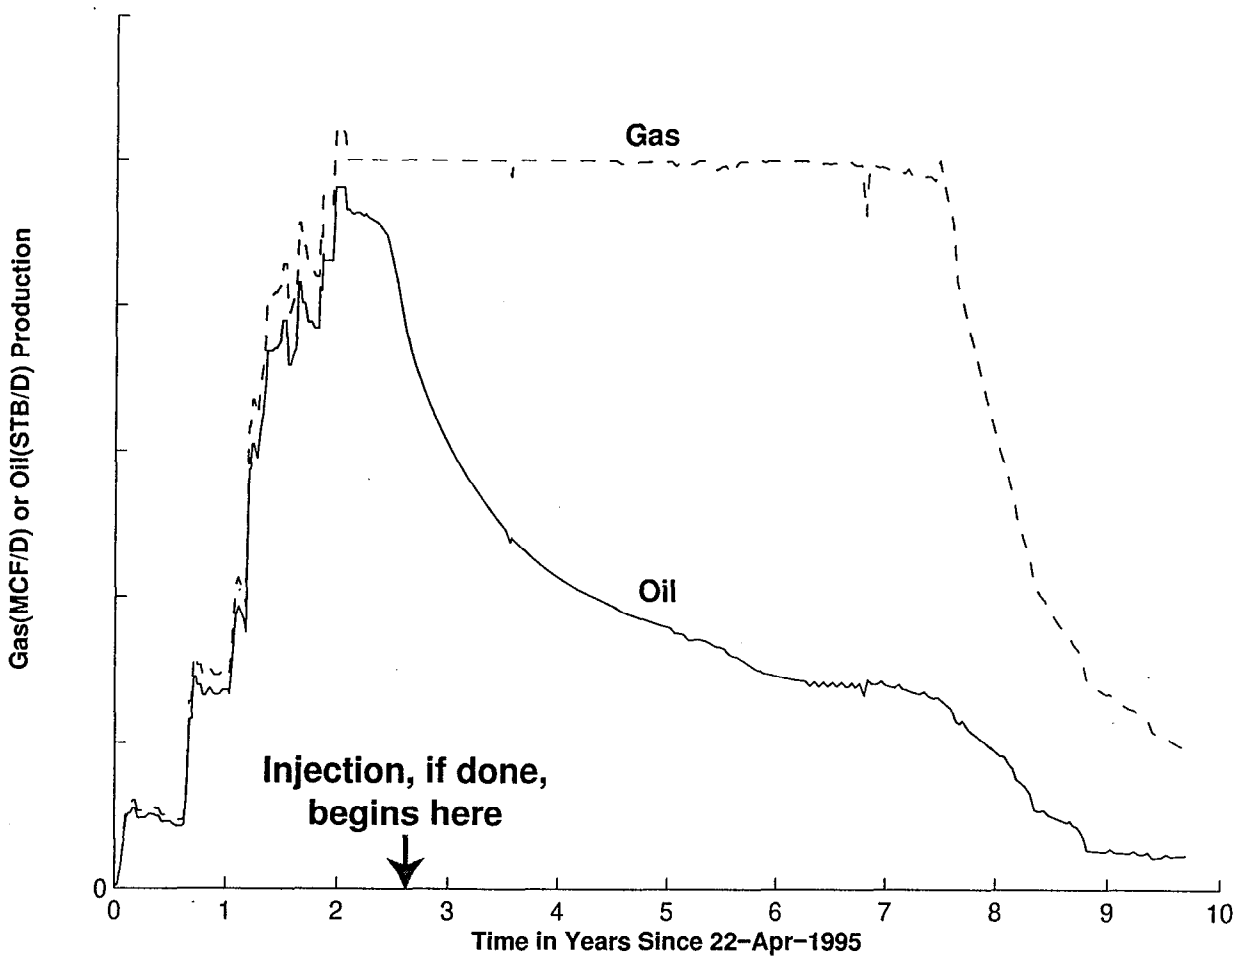

Fig. 4--Baseline Oil and Gas Production Decline Curves

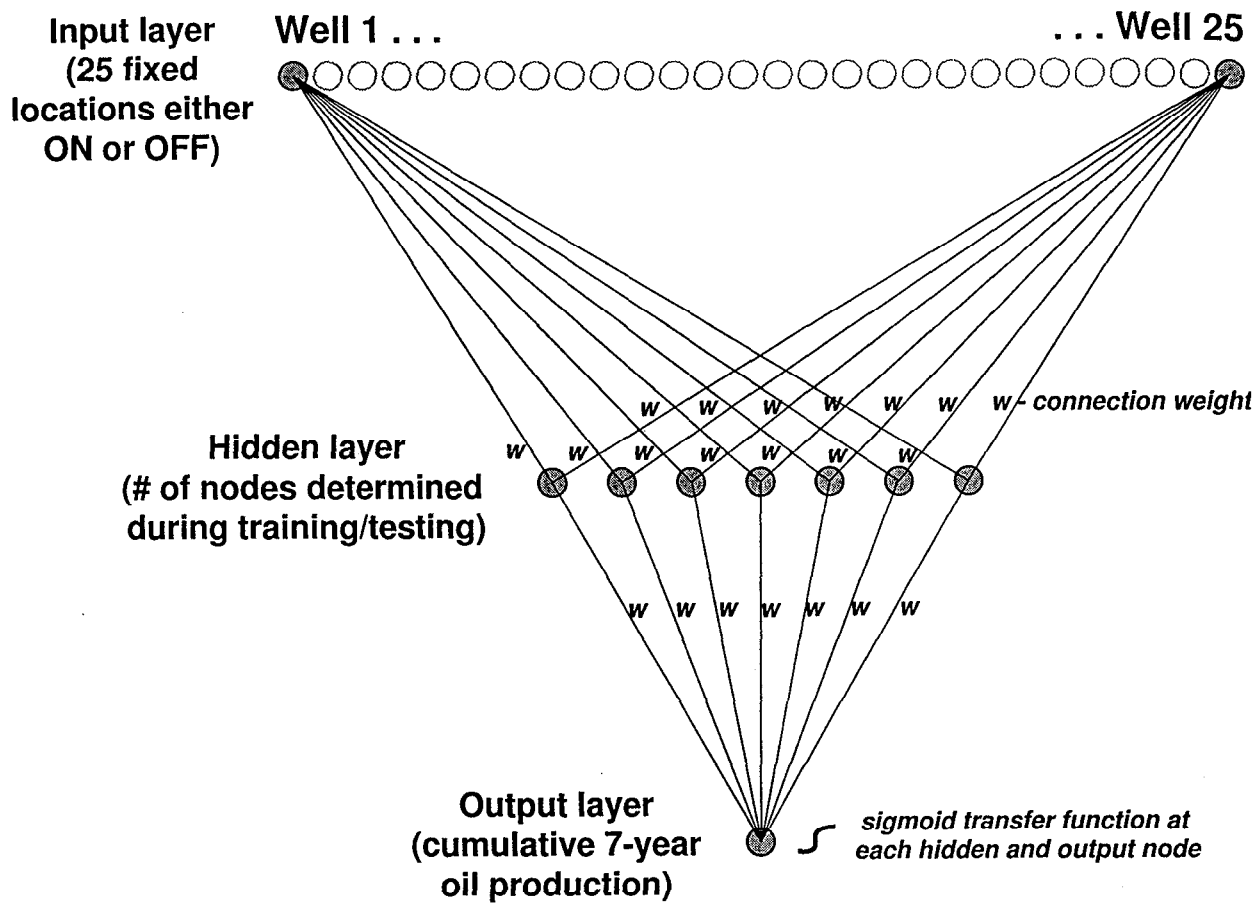

Fig. 5--Diagram of an ANN to predict cumulative 7-year oil production given a well combination as input. Connection weights between all input nodes except the first and last have been left out for visual simplicity.

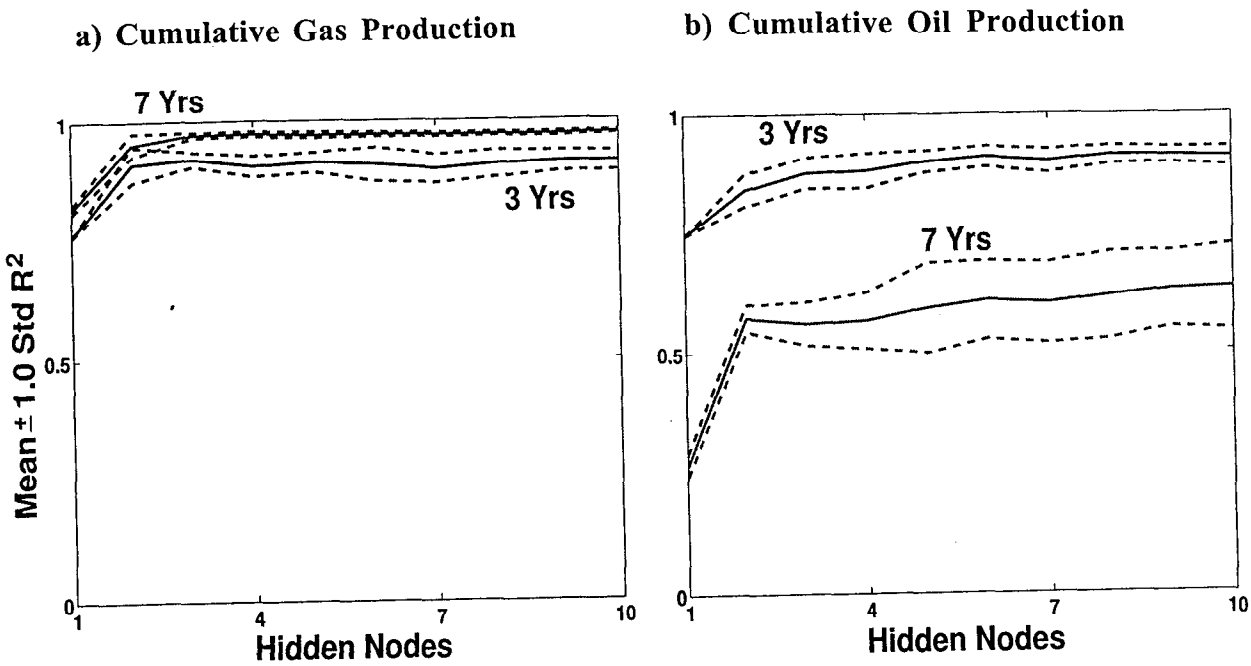

Fig. 6--Mean ANN generalization accuracy (shown by solid lines), ± one standard deviation (indicated by broken lines), as a function of the number of hidden nodes in the network, summarized over 25 random initializations of the network connection weights.